Knowledge Graph Reasoning and Its Applications

Lihui Liu
lihuil2@illinois.edu
University of Illinois at Urbana Champaign
Urbana, Illinois, USA

Hanghang Tong
htong@illinois.edu
University of Illinois at Urbana Champaign
Urbana, Illinois, USA

ABSTRACT

The use of knowledge graphs has gained significant traction in a wide variety of applications. By leveraging the wealth of information contained within knowledge graphs, it is possible to greatly enhance various downstream tasks. However, despite its popularity, knowledge graph reasoning remains a challenging problem. The first major challenge of knowledge graph reasoning lies in the nature of knowledge graphs themselves. Most knowledge graphs are incomplete, meaning that they may not capture all the relevant knowledge required for reasoning. As a result, reasoning on incomplete knowledge graphs can be difficult. Additionally, real-world knowledge graphs often evolve over time, which presents an additional challenge. The second challenge of knowledge graph reasoning pertains to the input data. In some KG reasoning applications, users may be unfamiliar with the background knowledge graph, leading to the possibility of asking ambiguous questions that can make KG reasoning tasks more challenging. Moreover, some applications require iterative reasoning, where users ask several related questions in sequence, further increasing the complexity of the task. The third challenge of knowledge graph reasoning concerns the algorithmic aspect. Due to the varied properties of relations in knowledge graphs, such as transitivity, symmetry, and asymmetry, designing an all-round KG reasoning model that fits all these properties can be challenging. This tutorial aims to comprehensively review different aspects of knowledge graph reasoning applications and highlight open challenges and future directions. It is intended to benefit researchers and practitioners in the fields of data mining, artificial intelligence, and social science. The slides can be found at

https://sites.google.com/view/kg-reasoning/home

CCS CONCEPTS

• **Computing methodologies** → **Reasoning about belief and knowledge**; • **Information systems** → **Data mining**.

KEYWORDS

Knowledge graph reasoning

ACM Reference Format:
Lihui Liu and Hanghang Tong. 2023. Knowledge Graph Reasoning and Its Applications. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '23), August 6–10, 2023, Long Beach, CA, USA.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

KDD '23, August 6–10, 2023, Long Beach, CA, USA

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM.
https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 AUDIENCE PARTICIPATION

The tutorial is aimed at researchers and practitioners in data mining, artificial intelligence, social science, and other interdisciplinary fields. Participants should have a basic understanding of probability, linear algebra, data mining, and machine learning, but no prior knowledge of knowledge graph reasoning is required. The tutorial is designed to accommodate participants with varying levels of expertise, with 40% of the material targeted at novices, 30% at intermediate learners, and 30% at experts, to ensure a good balance between introductory and advanced content. The tutorial will incorporate both lectures and audience discussions to encourage participation. Additionally, all materials, including tutorial descriptions, presentation slides, and pre-recorded videos, will be made available for post-tutorial dissemination.

2 PRESENTER BIOGRAPHY

Lihui Liu is a Ph.D. student in the Department of Computer Science at the University of Illinois at Urbana-Champaign. He is the corresponding tutor. His research focuses on large-scale data mining and machine learning, particularly on graphs, with an emphasis on knowledge graph reasoning. Lihui's research has been published at several major conferences and journals in data mining and artificial intelligence, and he has served as a reviewer and program committee member for top-tier data mining and artificial intelligence conferences and journals, including KDD, WWW, AAAI, IJCAI, and BigData. More information about Lihui can be found on his personal website at https://lihuiliullh.github.io/.

Hanghang Tong is currently an associate professor at Department of Computer Science at University of Illinois at Urbana-Champaign. Before that, he worked at Arizona State University as an associate professor, at City University of New York (City College) as an assistant professor and at IBM T. J. Watson Research Center as a Research Staff Member. He received his Ph.D. from the Machine Learning Department of School of Computer Science at Carnegie Mellon University in 2009. His major research interest lies in large-scale data mining for graphs and multimedia. In the past, He has published 200+ papers at these areas and his research has received several awards, including IEEE Fellow (2021), ACM distinguished member (2020), ICDM Tao Li award (2019), SDM/IBM Early Career Data Mining Research award (2018), NSF CAREER award (2017), ICDM 10-Year Highest Impact Paper award (2015), and several best paper awards. (e.g., ICDM'06 best paper, SDM'08 best paper, CIKM'12 best paper, etc.). He was Editor-in-Chief of ACM SIGKDD Explorations (2018 – 2022).

3 OUTLINE OF TUTORIAL

- **Introduction**
  We will discuss with the background knowledge for the topic, then briefly review existing problem definitions and settings, along with the key challenges in this topic. We will also introduce problems in relation to knowledge graph reasoning, including knowledge graph completion [1, 4, 10, 22, 23, 25], knowledge graph rule mining [2, 16, 18] ,knowledge graph question answering [5, 13, 19–21], knowledge graph fact checking [6, 12, 14, 15] and knowledge graph conversational question answering [3, 7, 8].
- **Part I: General Knowledge Graph Reasoning**
  In this part, we will discuss traditional knowledge graph reasoning, i.e., knowledge graph completion [1, 4, 9, 10, 22–25], rule mining [2, 16, 18] and others.
- **Part II: Query-specific Knowledge Graph Reasoning**
  In this part, we will discuss knowledge graph question answering [11, 13, 17, 19], and other applications, e.g., [3, 7, 8]
- **Part III: Summary and Future Directions**
  In this part, we summarize the tutorial and introduce potential future directions, such as recommendation, drug discovery and KG reasoning for large language models (LLMs)

4 RELATED TUTORIALS

Here, we discuss the most relevant tutorials in other conferences, as well as the similarities/differences compared with ours.

- **Reasoning on Knowledge Graphs: Symbolic or Neural?**
  -**Presenters**: Meng Qu, Zhaocheng Zhu, Jian Tang
  -**Conference**: AAAI 2022
  -**Connection**: This is an earlier version of this tutorial on knowledge graph reasoning
  -**Difference**: There will be some overlaps between the related tutorial and this tutorial (knowlege graph completion). However, in this tutorial, we will focus on knowledge graph reasoning applications, e.g., question answering, fact checking, instead of Symbolic Logic.

5 POTENTIAL SOCIETAL IMPACTS

This tutorial has several potential positive impacts to the society: (1) we hope this tutorial could attract research attention to promote fairness on graphs, which is less popular than fairness for IID data; (2) we hope this tutorial raises new challenges that are not addressed in the existing works; (3) we hope this tutorial could also benefit related research topic in identifying new problems and discovering its relationship to fairness on graphs

6 ACKNOWLEDGEMENT

This work is supported by NSF (1947135, and 2134079), the NSF Program on Fairness in AI in collaboration with Amazon (1939725), DARPA (HR001121C0165), NIFA (2020-67021-32799), DHS (17STQAC00001-06-00), ARO (W911NF2110088), the C3.ai Digital Transformation Institute, and IBM-Illinois Discovery Accelerator Institute. The content of the information in this document does not necessarily reflect the position or the policy of the Government or Amazon, and no official endorsement should be inferred. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

REFERENCES

[1] A Bordes and Usunier N. 2013. Translating Embeddings for Modeling Multi-relational Data. In *Advances in Neural Information Processing Systems 26.*
[2] Kewei Cheng, Jiahao Liu, Wei Wang, and Yizhou Sun. [n. d.]. RLogic: Recursive Logical Rule Learning from Knowledge Graphs. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '22).*
[3] Philipp Christmann and Rishiraj Saha Roy. 2019. Look before You Hop: Conversational Question Answering over Knowledge Graphs Using Judicious Context Expansion. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management (CIKM '19).*
[4] Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. 2017. Convolutional 2D Knowledge Graph Embeddings.
[5] W Hamilton, P Bajaj, M Zitnik, D Jurafsky, and J Leskovec. 2018. Embedding Logical Queries on Knowledge Graphs *(NIPS'18).* 2030–2041.
[6] Naeemul Hassan, Fatma Arslan, Chengkai Li, and Mark Tremayne. 2017. Toward Automated Fact-Checking: Detecting Check-Worthy Factual Claims by ClaimBuster. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.*
[7] Endri Kacupaj, Kuldeep Singh, Maria Maleshkova, and Jens Lehmann. 2022. Contrastive Representation Learning for Conversational Question Answering over Knowledge Graphs.
[8] Magdalena Kaiser, Rishiraj Saha Roy, and Gerhard Weikum. 2021. Reinforcement Learning from Reformulations in Conversational Question Answering over Knowledge Graphs. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '21).*
[9] Xi Victoria Lin, Richard Socher, and Caiming Xiong. 2018. Multi-Hop Knowledge Graph Reasoning with Reward Shaping. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing.*
[10] Y Lin, Z Liu, M Sun, Y Liu, and X Zhu. 2015. Learning Entity and Relation Embeddings for Knowledge Graph Completion. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence.*
[11] Lihui Liu, Yuzhong Chen, Mahashweta Das, Hao Yang, and Hanghang Tong. 2023. Knowledge Graph Question Answering with Ambiguous Query. In *Proceedings of the ACM Web Conference 2023.*
[12] Lihui Liu, Boxin Du, Yi Ren Fung, Heng Ji, Jiejun Xu, and Hanghang Tong. [n. d.]. KompaRe: A Knowledge Graph Comparative Reasoning System. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining.*
[13] Lihui Liu, Boxin Du, Jiejun Xu, Yinglong Xia, and Hanghang Tong. 2022. Joint Knowledge Graph Completion and Question Answering. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining.*
[14] Lihui Liu, Houxiang Ji, Jiejun Xu, and Hanghang Tong. 2022. Comparative Reasoning for Knowledge Graph Fact Checking. In *2022 IEEE International Conference on Big Data (Big Data).*
[15] Lihui Liu, Ruining Zhao, Boxin Du, Yi Ren Fung, Heng Ji, Jiejun Xu, and Hanghang Tong. 2022. Knowledge Graph Comparative Reasoning for Fact Checking: Problem Definition and Algorithms. *Data Engineering* (2022).
[16] Meng Qu, Junkun Chen, Louis-Pascal Xhonneux, Yoshua Bengio, and Jian Tang. 2021. RNNLogic: Learning Logic Rules for Reasoning on Knowledge Graphs. In *International Conference on Learning Representations.*
[17] Hongyu Ren, Weihua Hu, and Jure Leskovec. 2020. Query2box: Reasoning over Knowledge Graphs in Vector Space using Box Embeddings. In *International Conference on Learning Representations.*
[18] Ali Sadeghian, Mohammadreza Armandpour, Patrick Ding, and Daisy Zhe Wang. 2019. DRUM: End-To-End Differentiable Rule Mining On Knowledge Graphs.
[19] Apoorv Saxena and Aditay Tripathi. 2020. Improving multi-hop question answering over knowledge graphs using knowledge base embeddings. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics.*
[20] Haitian Sun and Tania Bedrax-Weiss. 2019. PullNet: Open Domain Question Answering with Iterative Retrieval on Knowledge Bases and Text.
[21] Haitian Sun, Bhuwan Dhingra, Manzil Zaheer, Kathryn Mazaitis, Ruslan Salakhutdinov, and William W. Cohen. 2018. Open Domain Question Answering Using Early Fusion of Knowledge Bases and Text. arXiv:1809.00782 [cs.CL]
[22] Z Sun, Z Deng, J Nie, and J Tang. [n. d.]. RotatE: Knowledge Graph Embedding by Relational Rotation in Complex Space. *CoRR* ([n. d.]).
[23] Théo Trouillon, Johannes Welbl, and Sebastian Riedel. [n. d.]. Complex Embeddings for Simple Link Prediction. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning (ICML'16).*
[24] Wenhan Xiong, Thien Hoang, and William Yang Wang. 2017. DeepPath: A Reinforcement Learning Method for Knowledge Graph Reasoning.
[25] Bishan Yang, Wen tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2015. Embedding Entities and Relations for Learning and Inference in Knowledge Bases.